arXiv:cs/0502004v1 [cs.LG] 1 Feb 2005

Asymptotic Log-loss of Prequential Maximum Likelihood Codes

Peter Grünwald and Steven de Rooij

February 1, 2008

Abstract

We analyze the Dawid-Rissanen *prequential maximum likelihood codes* relative to one-parameter exponential family models $\mathcal{M}$. If data are i.i.d. according to an (essentially) *arbitrary* P, then the redundancy grows at rate $\frac{1}{2}c \ln n$. We show that $c = \sigma_1^2/\sigma_2^2$, where σ_1^2 is the variance of P, and σ_2^2 is the variance of the distribution $M^* \in \mathcal{M}$ that is closest to P in KL divergence. This shows that prequential codes behave quite differently from other important universal codes such as the 2-part MDL, Shtarkov and Bayes codes, for which $c = 1$. This behavior is undesirable in an MDL model selection setting.

1 Introduction

Universal coding lies at the basis of on-line prediction algorithms for data compression and gambling purposes. It has been extensively studied in the COLT community, typically under the name of 'sequential prediction with log loss', see, for example [9, 1, 4]. It also underlies Rissanen's theory of MDL (minimum description length) learning [2, 11] and Dawid's theory of prequential model assessment [7]. Roughly, a code is *universal* with respect to a set of candidate codes $\mathcal{M}$ if it achieves small *redundancy*: it allows one to encode data using not many more bits than the optimal code in $\mathcal{M}$. The redundancy is very closely related to the *expected regret*. which is perhaps more widely known within the COLT community – we compare the two notions in Section 4. The main types of universal codes are the *Shtarkov* or *NML* code, the *Bayesian mixture* code, the *2-part MDL code* and the *prequential maximum likelihood (ML) code*, also known as the 'ML plug-in code' or the 'predictive MDL code' [2, 11]. This code was introduced independently by Rissanen [20] in the context of MDL learning and by Dawid [7], who proposed it as a probability forecasting strategy rather than directly as a code. The underlying ideas are explained in Section 2. Here we study the case where no code in $\mathcal{M}$ corresponds to the data generating distribution P. We find that in this case, the redundancy of the prequential code can be quite different from that of the other three methods. Specifically, if $\mathcal{M}$ is a one-dimensional exponential family, then the redundancies are $\frac{1}{2}c \ln n + O(1)$. Whereas it is known that for the Bayes, NML and 2-part codes, under regularity conditions on P and $\mathcal{M}$, we have $c = 1$ (Section 4), we determine c for the prequential code and find that, depending on properties of P and $\mathcal{M}$, it can be either larger or smaller than 1.

Relevance Our result has at least three important consequences, which are discussed further in Section 6.

1. **Practical consequence for data compression** When prequential codes are used for data compression in the realistic situation that $P \notin \mathcal{M}$, then depending on the situation they can behave either better or worse than the Bayesian and NML codes (end of Section 2).

2. **Practical consequence for MDL learning/model selection** In the case of model selection between two nonoverlapping parametric models, our results suggest (but do not prove) that the prequential plug-in codes typically behave worse (and never better) than the Bayesian or NML code. We have experimental evidence for this for the Poisson and geometric families.

3. **Theoretical** Our result implies that, under misspecification, the *Kullback-Leibler (KL) risk* of efficient estimators behaves in a fundamentally different way from the KL risk of estimators such as the Bayes predictive distribution which are not restricted to lie in the model $\mathcal{M}$ under consideration.

Contents The remainder of the paper is organized as follows. In Section 2 we informally state and explain our result, and we discuss how it relates to previous results. Section 3 contains the formal statement of our main result (Theorem 1), as well as a brief proof sketch. We show that a version of our result still holds if 'redundancy' is replaced by 'expected regret' in Section 4. We discuss further issues regarding our result in Section 5. We explain the relevance of our result, including the consequences listed above, in Section 6. Section 7 proves our main result. The proof makes use of several lemmas which are stated and proven in Section 8. The second result, discussed in Section 4, is proven in Section 9. The paper ends with a conclusion.

2 Main Result, Informally

Suppose $\mathcal{M} = \{M_\theta : \theta \in \Theta\}$ is a k-dimensional parametric family of distributions, and $Z_1, Z_2, \ldots$ are i.i.d. according to some distribution $P \in \mathcal{M}$. The *redundancy* of a universal code U with respect to P is defined as

$$\mathcal{R}_U(n) := E_P[L_U(Z_1, \ldots, Z_n)] - \inf_{\theta \in \Theta} E_P[-\ln M_\theta(Z_1, \ldots, Z_n)], \tag{1}$$

where L_U is the length function of U and $M_\theta(Z_1, \ldots, Z_n)$ denotes the probability mass or density of $Z_1, \ldots, Z_n$ under distribution M_θ; these and other notational conventions are detailed in Section 3. By the information inequality [6] the second term is minimized for $M_\theta = P$, so that

$$\mathcal{R}_U(n) = E_P[L_U(Z_1, \ldots, Z_n)] - E_P[-\ln P(Z_1, \ldots, Z_n)], \tag{2}$$

Thus, (2) can be interpreted as the expected number of additional nats one needs to encode n outcomes if one uses the code U instead of the optimal (Shannon-Fano) code with lengths $-\ln P(Z_1, \ldots, Z_n)$. A good universal code achieves small redundancy for all or 'most' $P \in \mathcal{M}$ (the relation to the concept of 'regret' is discussed in Section 4).

The four major types of universal codes, Bayes, NML, 2-part and prequential ML, all achieve redundancies that are (in an appropriate sense) close to optimal. Specifically, under regularity conditions on $\mathcal{M}$ and its parameterization, these four types of universal codes all satisfy, for all $P \in \mathcal{M}$,

$$\mathcal{R}(n) = \frac{k}{2} \ln n + O(1), \tag{3}$$

where the $O(1)$ may depend on θ and the universal code used. (3) is the famous 'k over 2 log n formula', refinements of which lie at the basis of most practical approximations to MDL learning, see [11].

In this paper we consider the case where the data are i.i.d. according to an *arbitrary* P not necessarily in the model $\mathcal{M}$. To emphasize that the redundancy is measured relative to the

element of the model that minimizes the codelength rather than to P, we use the term *relative redundancy* rather than just *redundancy*. Its definition (1) remains unchanged, but it can no longer be rewritten as (2): Assuming it exists and is unique, let M_{θ^*} be the element of $\mathcal{M}$ that minimizes KL divergence to P:

$$\theta^* := \arg\min_{\theta\in\Theta} D(P\|M_\theta) = \arg\min_{\theta\in\Theta} E_P[-\ln M_\theta(Z)],$$

where the equality follows from the definition of the KL divergence [6]. Then the relative redundancy satisfies

$$\mathcal{R}_U(n) = E_P[L_U(Z_1,\ldots,Z_n)] - E_P[-\ln M_{\theta^*}(Z_1,\ldots,Z_n)]. \tag{4}$$

It turns out that for the NML, 2-part MDL and Bayes codes, the relative redundancy (4) with $P \notin \mathcal{M}$, still satisfies (3), at least under conditions on $\mathcal{M}$ and P; see Section 4. In this paper, we show for the first time that (3) does *not* hold for the prequential ML code. The prequential ML code U works by sequentially predicting Z_{i+1} using a (slightly modified) ML or Bayesian MAP estimator $\hat\theta_i = \hat\theta(z^i)$ based on the past data, that is, the first i outcomes $z^i = z_1,\ldots,z_i$. The total codelength $L_U(z^n)$ on a sequence z^n is given by the sum of the individual 'predictive' codelengths (log losses): $L_U(z^n) = \sum_{i=0}^{n-1}[-\ln M_{\hat\theta_i}(z_{i+1})]$. In our main theorem, we show that if L_U denotes the prequential ML code length, and $\mathcal{M}$ is a regular one-parameter exponential family ($k=1$), then

$$\mathcal{R}_U(n) = \frac{1}{2}\frac{\mathrm{var}_P X}{\mathrm{var}_{M_{\theta^*}} X}\ln n + O(1), \tag{5}$$

where X is the sufficient statistic of the family. In Example 1 below we give an example of the phenomenon. The result holds as long as $\mathcal{M}$ and P satisfy Condition 1 defined below. Essentially, as long as the fourth moment of P exists, the condition holds for all exponential families we checked, including the Poisson, geometric, exponential, normal with fixed mean or variance and Pareto distributions. The result indicates that the redundancy can be both larger and smaller than $\frac{1}{2}\ln n$, depending on the variance of the 'true' P. We can only guarantee that the two variances are the same if $P \in \mathcal{M}$, in which case $M_{\theta^*} = P$. It immediately follows that in practical data compression tasks, whenever $P \notin \mathcal{M}$, the redundancy of the prequential ML code can be both smaller and larger than that of the Bayesian code, depending on the situation. This is the first of the three implications of our result, listed in Section 1. We postpone discussion of the other two implications to Section 6.

Example 1 Let $\mathcal{M}$ be the family of Poisson distributions, parameterized by their mean μ. Then the ML estimator $\hat\mu_i$ is the empirical mean of $z_1,\ldots,z_i$. Suppose Z, Z_1, Z_2, ... are i.i.d. according to a degenerate P with $P(Z=4)=1$. Since the sample average is a sufficient statistic for the Poisson family, $\hat\mu_i$ will be equal to 4 for all $i \geq 1$. On the other hand, μ^*, the parameter (mean) of the distribution in $\mathcal{M}$ closest to P in KL-divergence, will be equal to 4 as well. Thus the redundancy (4) of the prequential ML code is given by

$$\begin{aligned}\mathcal{R}_U(n) = \sum_{i=0}^{n-1}[-\ln M_{\hat\mu_i}(4) + \ln M_{\mu^*}(4)] &= -\ln M_{\hat\mu_0}(4) + \ln M_4(4) + \sum_{i=1}^{n-1}[-\ln M_4(4) + \ln M_4(4)] \\ &= -\ln M_{\hat\mu_0}(4) + \ln M_4(4) = O(1),\end{aligned} \tag{6}$$

assuming an appropriate definition of $\hat\mu_0$. In the case of the Poisson family, the outcome Z is equal to the sufficient statistic X in (5). Since $\mathrm{var}_P Z = 0$, this example agrees with (5).

Related Work There are a plethora of results concerning the redundancy and/or the regret for the prequential ML code, for a large variety of models including multivariate exponential families, ARMA processes, regression models and so on. Examples are [22, 10, 13, 25, 17]. In all these papers it is shown that either the regret or the redundancy grows as $\frac{k}{2}\ln n + o(\ln n)$, either in expectation or almost surely. [17] even evaluates the remainder term explicitly. The reason that these results do not contradict ours, is that in all these papers, one studies the case where the generating distribution P is in the model, in which case automatically $\mathrm{var}_{M^*}(X) = \mathrm{var}_P(X)$. In other cases [9, 21], regret of a prequential ML-type code is evaluated on an individual sequence basis, and it is found that the regret grows as $\frac{k}{2}\ln n + O(1)$ for all sequences whose ML estimator remains bounded away from the boundary of the space. The reason that these results do not contradict ours, is that in all cases that have been examined (and that we know of), the model is *complete*, i.e. it contains *all* distributions that can be defined on the sample space for 1 outcome. Then, if data are i.i.d. according to some P, P *must* be in $\mathcal{M}$, and we automatically get $\mathrm{var}_{M^*}(X) = \mathrm{var}_P(X)$. An example is [9] which uses the Bernoulli model. Apparently, we are the first to study the redundancy and regret for incomplete models under general circumstances.

3 Main Result, Formally

In this section, we introduce our notation, we define our quantities of interest, we state our main result and we give a short idea of the proof. The complete proof is given in Sections 7-8.

Notational Conventions Throughout this text we use nats rather than bits as units of information. Outcomes are capitalized if they are to be interpreted as random variables instead of instantiated values. A sequence of outcomes $z_1, \ldots, z_n$ is abbreviated to z^n. We write E_P as a shorthand for $E_{Z\sim P}$, the expectation of Z under distribution P. When we consider a sequence of n outcomes independently distributed $\sim P$, we use E_P even as a shorthand for the expectation of $(Z_1, \ldots, Z_n)$ under the n-fold product distribution of P. Finally, $P(Z)$ denotes the probability mass function of P in case Z is discrete-valued, and it denotes the density of P, in case Z takes its value in a continuum. When we write 'density function of Z', then, if Z is discrete-valued, this should be read as 'probability mass function of Z'. Note however that in our main result, Theorem 1 below, we do not assume that the data generating distribution P admits a density.

Exponential Families Let $\mathcal{Z}$ be a set of outcomes, taking values either in a finite or countable set, or in a subset of k-dimensional Euclidean space for some $k \geq 1$. Let $X : \mathcal{Z} \to \mathbb{R}$ be a random variable on $\mathcal{Z}$, and let $\mathcal{X} = \{x \in \mathbb{R} : \exists z \in \mathcal{Z} : X(z) = x\}$ be the range of X.

Exponential family models are families of distributions on $\mathcal{Z}$ defined relative to a random variable X (called 'sufficient statistic') as defined above, and a function $h : \mathcal{Z} \to [0, \infty)$. We let $\mathrm{Z}(\eta) := \int_{z\in\mathcal{Z}} e^{-\eta X(z)} h(z) dz$ (where the integral is to be replaced by a sum for countable $\mathcal{Z}$), and we let $\Theta_\eta := \{\eta \in \mathbb{R} : \mathrm{Z}(\eta) < \infty\}$.

Definition 1 (Exponential family) *The* single parameter exponential family [14] *with* sufficient statistic X *and* carrier h *is the family of distributions with densities* $M_\eta(z) := \frac{1}{\mathrm{Z}(\eta)} e^{-\eta X(z)} h(z)$, *where* $\eta \in \Theta_\eta$. Θ_η *is called the* natural parameter space. *The family is called* regular *if* Θ_η *is an open interval of* $\mathbb{R}$.

In the remainder of this text we only consider single parameter, regular exponential families where the mapping from Θ_η to the corresponding set of distributions is 1-to-1, but these qualifications will henceforth be omitted. Examples of this wide family of models include the Poisson,

geometric and multinomial families, and the model of all Gaussian (normal) distributions with a fixed variance, or with a fixed mean. In the first four cases, we can take X to be the identity, so that $X = Z$ and $\mathcal{X} = \mathcal{Z}$. In the case of the normal family with fixed mean, σ^2 becomes the sufficient statistic and we have $\mathcal{Z} = \mathbb{R}$, $\mathcal{X} = [0, \infty)$ and $X = Z^2$.

The statistic $X(z)$ is sufficient for η [14]. This suggests reparameterizing the distribution by the expected value of X, which is called the *mean value parameterization*. The function $\mu(\eta) = E_{M_\eta}[X]$ maps parameters in the natural parameterization to the mean value parameterization. It is a diffeomorphism (it is one-to-one, onto, infinitely often differentiable and has an infinitely often differentiable inverse) [14]. Therefore the mean value parameter space Θ_μ is also an open interval of $\mathbb{R}$. We note that for some models (such as Bernoulli and Poisson), the parameter space is usually given in terms of the a non-open set of mean-values (e.g., $[0, 1]$ in the Bernoulli case). In this case, to make the model a regular exponential family, we have to restrict the set of parameters to its own interior. Henceforth, whenever we refer to a standard statistical model such as Bernoulli or Poisson, we assume that the parameter set has been restricted in this sense.

We are now ready to define the prequential ML model. This is a distribution on infinite sequences $z_1, z_2, \ldots \in \mathcal{Z}^\infty$, recursively defined in terms of the distributions of Z_{n+1} conditioned on $Z^n = z^n$, for all $n = 1, 2, \ldots$, all $z^n = (z_1, \ldots, z_n) \in \mathcal{Z}^n$. In the definition, we use the notation $x_i := X(z_i)$.

Definition 2 (Prequential ML model) *Let Θ_μ be the mean value parameter domain of an exponential family $\mathcal{M} = \{M_\mu \mid \mu \in \Theta_\mu\}$. Given $\mathcal{M}$ and constants $x_0 \in \Theta_\mu$ and $n_0 > 0$, we define the* prequential ML model *U by setting, for all n, all $z^{n+1} \in \mathcal{Z}^{n+1}$:*

$$U(z_{n+1} \mid z^n) = M_{\hat{\mu}(z^n)}(z_{n+1}),$$

where $U(z_{n+1} \mid z^n)$ is the density/mass function of z_{n+1} conditional on $Z^n = z^n$,

$$\hat{\mu}(z^n) := \frac{x_0 \cdot n_0 + \sum_{i=1}^n x_i}{n + n_0},$$

and $M_{\hat{\mu}(z^n)}(\cdot)$ is the density of the distribution in $\mathcal{M}$ with mean $\hat{\mu}(z^n)$.

We henceforth abbreviate $\hat{\mu}(z^n)$ to $\hat{\mu}_n$. We usually refer to the prequential ML model in terms of the corresponding codelength function

$$L_U(z^n) = \sum_{i=0}^{n-1} L_U(z_{i+1} \mid z_i) = \sum_{i=0}^{n-1} -\ln M_{\hat{\mu}_i}(z_{i+1}).$$

To understand this definition, note that for exponential families, for any sequence of data, the ordinary maximum likelihood parameter is given by the average $n^{-1} \sum x_i$ of the observed values of X [14]. Here we define our prequential model in terms of a slightly modified maximum likelihood estimator that introduces a 'fake initial outcome' x_0 with multiplicity n_0 in order to avoid infinite code lengths (see the quote by Rissanen on "inherent singularity" in Section 6) and to ensure that the prequential ML code length of the first outcome is well-defined. In practice we can take $n_0 = 1$ but our result holds for any $n_0 > 0$. This definition can be reconciled with settings in which the startup problem is resolved by ignoring the first few outcomes, by setting x_0 to the ML estimator for the ignored outcomes and n_0 to their number. It also allows our results to be generalized to a number of other point estimators as discussed in Section 5.2.

With all our definitions in place we can state our main result.

Theorem 1 (Main result) *Let $X, X_1, X_2, \ldots$ be i.i.d. $\sim P$, with $E_P[X] = \mu^*$. Let $\mathcal{M}$ be a single parameter exponential family with sufficient statistic X and μ^* an element of the mean value parameter space. Finally let U denote the prequential ML model with respect to $\mathcal{M}$. If $\mathcal{M}$ and P satisfy Condition 1 below, then*

$$\mathcal{R}_U(n) = \frac{\mathrm{var}_P X}{\mathrm{var}_{M_{\mu^*}} X} \frac{1}{2} \ln n + O(1).$$

To reconcile this with the informal statement (5), notice that M_{μ^*} is the element of $\mathcal{M}$ achieving the smallest expected codelength, i.e. it achieves $\inf_{\mu \in \Theta_\mu} D(P \| M_\mu)$ [14].

Condition 1 *We require that the following holds both for $T := X$ and $T := -X$:*

- *If T is unbounded from above then there is a $k \in \{4, 6, \ldots\}$ such that the first k moments of T exist under P and that $\frac{d^4}{d\mu^4} D(M_{\mu^*} \| M_\mu) = O\left(\mu^{k-6}\right)$.*
- *If T is bounded from above by a constant g then $\frac{d^4}{d\mu^4} D(M_{\mu^*} \| M_\mu)$ is polynomial in $1/(g-\mu)$.*

The condition implies that Theorem 1 can be applied to most single-parameter exponential families that are relevant in practice. To illustrate, we have computed the fourth derivative of the divergence for a number of exponential families; all parameters beside the mean are treated as fixed values. The results are listed in Figure 1. As can be seen from the figure, for these exponential families, our condition applies whenever the fourth moment of P exists. Note in particular that the condition requires $\mathrm{var}_P X < \infty$.

The reason why we need Condition 1 is best explained by sketching the proof of Theorem 1:

Brief Proof Sketch The precise proof of Theorem 1, given in Section 7, is very technical. Here we merely describe the underlying ideas, which are relatively simple. Consider first the case $n_0 = 0$, so that for $n \geq 1$, $\hat{\mu}_n$ is just the standard ML estimator. Let z^i be the initial i outcomes of an arbitrary sequence $z^n = z_1, z_2, \ldots, z_n$. As is well-known, a straightforward second-order Taylor expansion of $D(M_{\mu^*} \| M_{\hat{\mu}(z^i)})$ around μ^* gives

$$D(M_{\mu^*} \| M_{\hat{\mu}_i}) = \frac{1}{2} I(\mu^*)(\hat{\mu}_i - \mu^*)^2 + \text{Remainder}. \tag{7}$$

Here $I(\mu^*)$ is the Fisher information in one observation, evaluated at μ^*, see Section 7. For exponential families in their mean-value parameterization, another standard result [14] says that for all μ,

$$I(\mu) = \frac{1}{\mathrm{var}_{M_\mu} X} \tag{8}$$

Therefore, ignoring the remainder term and the term for $i = 0$, we get

$$\sum_{i=0}^{n-1} \underset{\hat{\mu}_i \sim P}{E} [D(M_{\mu^*} \| M_{\hat{\mu}_i})] \approx \frac{1}{2} \sum_{i=1}^{n-1} \frac{E_P(\hat{\mu}_i - \mu^*)^2}{\mathrm{var}_{M_{\mu^*}} X} =$$

$$\frac{1}{2} \frac{\mathrm{var}_P X}{\mathrm{var}_{M_{\mu^*}} X} \sum_{i=1}^{n-1} \frac{1}{i} = \frac{1}{2} \frac{\mathrm{var}_P X}{\mathrm{var}_{M_{\mu^*}} X} \ln n + O(1), \tag{9}$$

Here the first approximate equality follows by (7) and (8). The second follows because for exponential families, the ML estimator $\hat{\mu}_i$ is just the empirical average $i^{-1} \sum x_i$, so that $E_P(\hat{\mu}_i -$

	$M_{\mu^*}(x)$	$\frac{d^4}{d\mu^4} D(M_{\mu^*} \| M_\mu)$
Bernoulli	$(\mu^*)^x (1-\mu^*)^{(1-x)}$	$\frac{6\mu^*}{\mu^4} + \frac{6(1-\mu^*)}{(1-\mu)^4}$
Poisson	$\frac{e^{\mu^*} \mu^{*x}}{x!}$	$\frac{6\mu^*}{\mu^4}$
Geometric	$\theta^x(1-\theta) = \frac{(\mu^*)^x}{(\mu^*+1)^{x+1}}$	$\frac{6\mu^*}{\mu^4} - \frac{6(\mu^*+1)}{(\mu+1)^4}$
Exponential	$\frac{1}{\mu^*} e^{-x/\mu^*}$	$-\frac{6}{\mu^4} + \frac{24\mu^*}{\mu^5}$
Normal (fixed mean = 0)	$\frac{1}{\sqrt{2\pi\mu^* x}} e^{-\frac{x}{2\mu^*}}$	$-\frac{3}{\mu^4} + \frac{12\mu^*}{\mu^5}$
Normal (fixed variance = 1)	$\frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(x-\mu^*)^2}$	0
Pareto	$\frac{ab^a}{x^{a+1}}$ for $b = \frac{a-1}{a}\mu^*$	$\frac{6a}{\mu^4}$

Figure 1: $\frac{d^4}{d\mu^4} D(M_{\mu^*} \| M_\mu)$ for a number of exponential families. For the normal distribution we use mean 0, and we list a reparametrization of the density function such that the density of the squared outcomes is given as a function of the variance, which is confusingly but correctly called μ^* here: the random variable X in Theorem 1 is really the observed value of z^2 rather than z itself, so that its mean is $E[X] = E[Z^2]$, which is the variance of the normal distribution.

$\mu^*)^2 = \mathrm{var}(i^{-1} \sum_{j=1}^{i} X_j) = i^{-1} \mathrm{var} X$. Thus, Theorem 1 follows if we can show (a) that the left-hand side of (9) is equal to the relative redundancy $\mathcal{R}_U(n)$ and (b) that, as $n \to \infty$, the remainder terms in (7), summed over n as in (9), form a convergent series (i.e. sum to something finite). Result (a) follows relatively easily by rewriting the sum using the chain rule for relative entropy and using the fact that X is a sufficient statistic (Lemma 9). The truly difficult part of the proof is (b), shown in Lemma 11. It involves infinite sums of expectations over unbounded fourth-order derivatives of the KL divergence. To make this work, we (1) slightly modify the ML estimator by introducing the initial fake outcome x_0. And (2), we need to impose Condition 1. To understand it, consider the case $T = X$, X unbounded from above. The condition essentially expresses that, as $\hat{\mu}$ increases to infinity, the fourth order Taylor-term does not grow too fast. Similarly, if X is bounded from above by g, the condition ensures that the fourth-order term grows slowly enough as $\hat{\mu} \uparrow g$. The same requirements are imposed for decreasing $\hat{\mu}$.

4 Redundancy vs. Regret

The 'goodness' of a universal code relative to a model $\mathcal{M}$ can be measured in several ways: rather than using redundancy (as we did here), one can also choose to measure codelength differences in terms of *regret*, where one has a further choice between *expected regret* and *worst-case regret* [2]. Here we only discuss the implications of our result for the expected regret measure.

Let $\mathcal{M} = \{M_\theta \mid \theta \in \Theta\}$ be a family of distributions parameterized by Θ. Given a sequence $z^n = z_1, \ldots, z_n$ and a universal code U for $\mathcal{M}$ with lengths L_U, the *regret* of U on sequence z^n

is defined as

$$L_U(z^n) - \inf_{\theta\in\Theta}[-\ln M_\theta(z^n)]. \tag{10}$$

Note that if the (unmodified) ML estimator $\hat{\theta}(z^n)$ exists, then this is equal to $L_U(z^n) + \ln M_{\hat{\theta}(z^n)}(z^n)$. Thus, one compares the codelength achieved on z^n by U to the best possible that could have been achieved on that particular z^n, using any of the distributions in $\mathcal{M}$. Assuming $Z_1, Z_2, \ldots$ are i.i.d. according to some (arbitrary) P, one may now consider the expected regret

$$\widehat{\mathcal{R}}_U(n) := E_P[L_U(Z^n) - \inf_{\theta\in\Theta}[-\ln M_\theta(Z^n)]],$$

To quantify the difference between redundancy and expected regret, consider the function

$$d(n) := \inf_{\theta\in\Theta} E_P[-\ln M_\theta(Z^n)] - E_P\,[\inf_{\theta\in\Theta}[-\ln M_\theta(Z^n)]],$$

and note that for any universal code, $\mathcal{R}_U(n) - \widehat{\mathcal{R}}_U(n) = d(n)$. In case $P \in \mathcal{M}$, then under regularity conditions on $\mathcal{M}$ and its parameterization, it can be shown [5] that

$$\lim_{n\to\infty} d(n) = \frac{k}{2}, \tag{11}$$

where k is the dimension of $\mathcal{M}$. In our case, where P is not necessarily in $\mathcal{M}$, we have the following:

Theorem 2 *Let $\mathcal{X}$ be finite. Let P, M_μ and μ^* be as in Theorem 1. Then*

$$\lim_{n\to\infty} d(n) = \frac{1}{2}\frac{\mathrm{var}_P X}{\mathrm{var}_{M_{\mu^*}} X}. \tag{12}$$

Once we are dealing with 1-parameter families, in the special case that $P \in \mathcal{M}$, this result reduces to (11). We conjecture that, under a condition similar to Condition 1, the same result still holds for general, not necessarily finite or countable or bounded $\mathcal{X}$, but at the time of writing this submission we did not yet find the time to sort out the details. In any case, our result is sufficient to show that in some cases (namely, if $\mathcal{X}$ is finite), we have

$$\widehat{\mathcal{R}}_U(n) = \frac{1}{2}\frac{\mathrm{var}_P X}{\mathrm{var}_{M_{\mu^*}} X}\ln n + O(1),$$

so that, up to $O(1)$-terms, the redundancy and the regret of the prequential ML code behave in the same way.

Incidentally, Theorem 2 can be used to substantiate the claim we made in Section 2, which stated that the Bayes (equipped with a strictly positive differentiable prior), NML and 2-part codes still achieve relative redundancy of $\frac{1}{2}\ln n$ if $P \neq \mathcal{M}$, at least if $\mathcal{X}$ is finite. Let us informally explain why this is the case. It is easy to show that Bayes, NML and (suitably defined) 2-part codes achieve regret $\frac{1}{2}\ln n + O(1)$ for *all* sequences $z_1, z_2, \ldots$ such that $\hat{\theta}(z^n)$ is bounded away from the boundary of the parameter space M, for all large n [2, 11]. It then follows using, for example, the Chernoff bound that these codes must also achieve expected regret $\frac{1}{2}\ln n + O(1)$ for *all* distributions P on $\mathcal{X}$ that satisfy $E_P[X] = \mu^* \in \Theta_\mu$. Theorem 2 then shows that they also achieve relative redundancy $\frac{1}{2}\ln n + O(1)$ for *all* distributions P on $\mathcal{X}$ that satisfy $E_P[X] = \mu^* \in \Theta_\mu$. We omit further details.

5 Variations of Prequential Coding

5.1 Justifying Our Modification of the ML Estimator

If the prequential code is based on the ordinary ML estimator ($n_0 = 0$ in Definition 2) then, apart from being undefined for the first outcome, it may achieve infinite codelengths on the observed data. A simple example is the Bernoulli model. If we first observe $z_1 = 0$ and then $z_2 = 1$, the codelength of z_2 according to the ordinary ML estimator of z_2 given z_1 would be $-\ln M_{\hat{\mu}}(z_1)(z_2) = -\ln 0 = \infty$. There are several ways to resolve this problem. We choose to add an 'initial fake outcome'. Another possibility that has been suggested (e.g., [7]) is to use the ordinary ML estimator, but only start using after having observed m examples, where m is the smallest number such that $-\ln M_{\hat{\mu}(z^m)}(Z_{m+1})$ is guaranteed to be finite, no matter what value Z_{m+1} is realized. The first m outcomes may then be encoded by repeatedly using some code L_0 on outcomes of $\mathcal{Z}$, so that for $i \leq m$, the codelength of z_i does not depend on the outcomes z^{i-1}. In the Bernoulli example, one could for example use the code corresponding to $P(Z_i = 1) = 1/2$, until and including the first i such that z^i includes both a 0 and a 1. Then it takes i bits to encode the first z^i outcomes, no matter what they are. After that, one uses the prequential code with the standard ML estimator. It is easy to see (by slight modification of the proof) that our theorem still holds for this variation of prequential coding. Thus, our particular choice for resolving the startup problem is not crucial to obtain our result. The advantage of our solution is that, as we now show, it allows us to interpret our modified ML estimator also as a Bayesian MAP and Bayesian mean estimator, thereby showing that the same behavior can be expected for such estimators.

5.2 Prequential Models with Other Estimators

The Bayesian MAP estimator If a conjugate prior is used, the Bayesian maximum a-posteriori estimator can always be interpreted as an ML estimator based on the sample and some additional 'fake data' ([3]; see also the notion of *ESS (Equivalent Sample Size) Priors* discussed in, for example, [15]). Therefore, the prequential ML model as defined above can also be interpreted as a prequential MAP model for that class of priors, and the whole analysis carries over to that setting.

The Bayesian mean estimator It follows by the work of Hartigan [12, Chapter 7] on the so-called 'maximum likelihood prior', that by slightly modifying conjugate priors, we can construct priors such that the Bayesian *mean* rather than MAP estimator is of the form of our modified ML estimator.

A Conjecture In some special cases, for example, the Bernoulli model, the exponential family $\mathcal{M}$ covers all distributions that can be defined on $\mathcal{X}$. In such cases, there exists no distribution with mean μ^* and variance not equal to $\mathrm{var}_{M_{\mu^*}} X$, and the $\frac{1}{2}\ln n + O(1)$ redundancy can always be achieved. But in all other cases, three very reasonable and efficient [19] estimators (ML, Bayes MAP, Bayes mean for a large class of reasonable priors) cannot achieve $\frac{1}{2}\ln n + O(1)$ in all circumstances. This suggests that *no matter what in-model*[1] *estimator is used*, the prequential model cannot yield a relative redundancy of $\frac{1}{2}\ln n$ independently of the variance of the data generating distribution P.

[1] See Section 6.

5.3 Rissanen's Predictive MDL Approach

The MDL model selection criterion that is based on comparing the prequential ML codelengths for the models under consideration is called the *Predictive MDL (PMDL)* criterion by Rissanen [22]. It is closely related to the *Predictive Least Squares (PLS)* criterion [24] for regression models; PMDL can be seen as an MDL justification for it. There has been some discussion on how to use PMDL when the data are not ordered. The prequential ML codelength then becomes redundant: the same data can be coded in any order, yielding different code words. Rissanen suggests in [23] to use the permutation of the outcomes that minimizes the codelength. Under such a regime, Theorem 1 is no longer applicable (since the outcomes are no longer i.i.d.); however Example 1 illustrates that circumstances in which the prequential ML codelength and the NML codelength behave very differently remain, under *any* regime that amounts to reordering the sample, including the one suggested by Rissanen.

6 Consequences

Why are these results interesting? We listed three significant implications in Section 1, the introduction to this paper. The first was evident from Theorem 1. Let us now discuss the second and third in more detail.

Practical significance for Model Selection There exist a plethora of results showing that in various contexts, if $P \in \mathcal{M}$, then the prequential ML code achieves optimal redundancy (see Section 2, Related Work). These strongly suggest that it is a very good alternative for (or at least approximation to) the NML or Bayesian codes in MDL model selection. Indeed, quoting Rissanen [24]:

> "If the encoder does not look ahead but instead computes the best parameter values from the past string, only, using an algorithm which the decoder knows, then no preamble is needed. The result is a *predictive* coding process, one which is quite different from the sum or integral formula in the stochastic complexity.[2] And it is only because of a certain inherent singularity in the process, as well as the somewhat restrictive requirement that the data must be ordered, that we do not consider the resulting predictive code length to provide another competing definition for the stochastic complexity, but rather regard it as an approximation."

Our result however shows that the prequential ML code may behave quite differently from the NML and Bayes codes, thereby strengthening the conclusion that it should not be taken as a definition of stochastic complexity. Although there is only a significant difference if data are distributed according to some $P \notin \mathcal{M}$, the difference is nevertheless very relevant in an MDL model selection context with nonoverlapping models, even if one of the models under consideration *does* contain the 'true' P. To see this, suppose we are comparing two models $\mathcal{M}_1$ and $\mathcal{M}_2$ for the same data, and in fact, $P \in \mathcal{M}_1 \cup \mathcal{M}_2$. For concreteness, assume $\mathcal{M}_1$ is the Poisson family and $\mathcal{M}_2$ is the geometric family. We want to decide which of these two models best explains the data. According to the MDL Principle, we should associate with each model a universal code (preferably the NML code). We should then pick the model such that the corresponding universal codelength of the data is minimized. Now suppose we use the prequential ML codelengths rather than the NML codelengths. Without loss of generality suppose that $P \in \mathcal{M}_1$. Then $P \notin \mathcal{M}_2$. This means that the codelength relative to $\mathcal{M}_1$ behaves

[2] The stochastic complexity is the codelength of the data $z_1, \ldots, z_n$ that can be achieved using the NML code.

essentially like the NML codelength, but the codelength relative to $\mathcal{M}_2$ behaves differently – at least as long as the variances do not match (which for example, is forcibly the case if $\mathcal{M}_1$ is Poisson and $\mathcal{M}_2$ is geometric). This introduces a bias in the model selection scheme. We have found experimentally [8] that the error rate for model selection based on the prequential ML code decreases more slowly than when other universal codes are used. Even though in some cases the redundancy grows *more slowly* than $\frac{1}{2}\ln n$, so that the prequential ML code is in a sense more efficient than the NML code, model selection based on the prequential ML codes behaves worse than Bayesian and NML-based model selection. We provide a theoretical explanation for this phenomenon in [8]. The practical relevance of this phenomenon stems from the fact that the prequential ML codelengths are often a lot easier to compute than the Bayes or NML codes, so that they are often used in applications [18, 16].

Theoretical Significance The result is also of theoretical-statistical interest: our theorem can be re-interpreted as establishing bounds on the asymptotic *Kullback-Leibler risk* of density estimation using ML and Bayes estimators under misspecification ($P \notin \mathcal{M}$). Our result implies that, under misspecification, the KL risk of estimators such as ML, which are required to lie in the model $\mathcal{M}$, behaves in a fundamentally different way from the KL risk of estimators such as the Bayes predictive distribution, which are not restricted to lie in $\mathcal{M}$. Namely, we can think of every universal model U defined as a random process on infinite sequences as an *estimator* in the following way: define, for all n,

$$\breve{P}_n := \Pr_U(Z_{n+1} = \cdot \mid Z_1 = z_1, \ldots, Z_n = z_n),$$

a function of the sample $z_1, \ldots, z_n$. $\breve{P}_n$ can be thought of as the 'estimate of the true data generating distribution upon observing $z_1, \ldots, z_n$'. In case U is the prequential ML model, $\breve{P}_n = M_{\hat{\theta}_n}$ is simply our modified ML estimator. It is now important to note that for other universal models, $\breve{P}_n$ is not required to lie in $\mathcal{M}$. An example is the Bayesian universal code defined relative to some prior w. This code has lengths $L'(z^n) := -\ln \int M_\mu(z^n) w(\mu) d\mu$ [11]. The corresponding estimator is the *Bayesian posterior predictive distribution* $P_{\mathrm{Bayes}}(z_{i+1} \mid z^i) := \int M_\mu(z_{i+1}) w(\mu \mid z^i) d\mu$ [11]. The Bayesian predictive distribution is a mixture of elements of $\mathcal{M}$. We will call standard estimators like the ML estimator, which are required to lie in $\mathcal{M}$, *in-model* estimators. Estimators like the Bayesian predictive distribution will be called *out-model*.

Let now $\breve{P}_n$ be any estimator, in-model or out-model. Let $\breve{P}_{z^n}$ be the distribution estimated for a particular realized sample z^n. We can measure the closeness of $\breve{P}_{z^n}$ to M_{μ^*}, the distribution in $\mathcal{M}$ closest to P in KL-divergence, by considering the *extended KL divergence*

$$D^*(M_{\mu^*} \| \breve{P}_{z^n}) = E_{Z \sim P}[-\ln \breve{P}_{z^n}(Z) - [-\ln M_{\mu^*}(Z)]].$$

We can now consider the *expected* KL divergence between M_{μ^*} and $\breve{P}_n$ after observing a sample of length n:

$$E_{Z_1,\ldots,Z_n \sim P}[D^*(M_{\mu^*} \| \breve{P}_n)]. \tag{13}$$

In analogy to the definition of 'ordinary' *KL risk* [2], we call (13) the *extended KL risk*. We recognize $\mathcal{R}_U(n)$, the redundancy of the prequential ML model, as the accumulated expected KL risk of our modified ML estimator (see Proposition 10 and Lemma 9). In exactly the same way as for the prequential ML code, the redundancy of the Bayesian code can be re-interpreted as the accumulated KL risk of the Bayesian predictive distribution. With this interpretation, our Theorem 1 expresses that under misspecification, the cumulative KL risk of the ML estimator differs from the cumulative KL risk of the Bayes estimator by a term of $O(\ln n)$. If our conjecture that *no* in-model estimator can achieve redundancy $\frac{1}{2}\ln n + O(1)$ for all μ^* and all P with finite

variance is true (Section 5.2), then it follows that the KL risk for in-model estimators behaves in a fundamentally different way from the KL risk for out-model estimators, and that out-model estimators are needed to achieve the optimal constant $c = 1$ in the redundancy $\frac{1}{2}c \ln n + O(1)$.

7 Proof of Theorem 1

Preliminaries Note that, for any $M_\mu, M_{\mu'} \in \mathcal{M}$, we have

$$\begin{aligned} E_P[-\ln M_\mu(Z)] - E_P[-\ln M_{\mu'}(Z)] &= \eta(\mu)E_P[X(Z)] + \ln \mathrm{Z}(\eta(\mu)) + E_P[-\ln h(Z)] \\ &- \eta(\mu')E_P[X(Z)] - \ln \mathrm{Z}(\eta(\mu')) - E_P[-\ln h(Z)] \\ &= E_P[-\ln M_\mu(X)] - E_P[-\ln M_{\mu'}(X)], \end{aligned}$$

so that we have

Proposition 3

$$\mathcal{R}_U(n) = E_P[-\ln M(X^n)] - \inf_\mu E_P[-\ln M_\mu(X^n)].$$

Proposition 3 shows that relative redundancy, which is the sole quantity of interest in the proof, depends only on the value of X, not Z. Thus, in the proof of Theorem 1 as well as all the Lemmas and Propositions it makes use of, we will never mention Z again. Whenever we refer to a 'distribution' we mean a distribution of random variable X, and we also think of the data generating distribution P in terms of the distribution it induces on X rather than Z. Whenever we say 'the mean' without further qualification, we refer to the mean of the random variable X. Whenever we refer to the Kullback-Leibler (KL) divergence between P and Q, we refer to the KL divergence between the distributions they induce for X (the reader who is confused by this may simply restrict attention to exponential family models for which $Z = X$, and consider X and Z identical).

The proof refers to a number of theorems and lemmas which will be developed in Section 8. In the statement of all these results, we assume, as in the statement of Theorem 1, that $X, X_1, X_2, \ldots$ are i.i.d. $\sim P$ and that μ^* is the mean of X under P. If X takes its values in a countable set, then all integrals in the proof should be read as the corresponding sums.

Proof (of Theorem 1) From Lemma 9 we have:

$$\mathcal{R}_U(n) = \sum_{i=0}^{n-1} \underset{\hat{\mu}_i \sim P}{E} [D(M_{\mu^*} \| M_{\hat{\mu}_i})] \tag{14}$$

Here, $\hat{\mu}_i$ is a random variable that takes on values according to P, while μ^* is fixed. We first abbreviate $\delta_i = \hat{\mu}_i - \mu^*$ and $\frac{d^k}{d\mu^k} D(M_{\mu^*} \| M_\mu) = D^{(k)}(\mu)$, That is, $D^{(k)}(\mu)$ is the k-th derivative of the function $f(\mu) := D(M_{\mu^*} \| M_\mu)$. We now Taylor-expand the divergence around μ^*:

$$D(M_{\mu^*} \| M_{\hat{\mu}_i}) = 0 + \delta_i D^{(1)}(\mu^*) + \frac{{\delta_i}^2}{2} D^{(2)}(\mu^*) + \frac{{\delta_i}^3}{6} D^{(3)}(\mu^*) + \frac{{\delta_i}^4}{24} D^{(4)}(\mu)$$

The last term is the remainder term of the Taylor expansion, in which $\mu \in [\mu^*, \hat{\mu}_i]$. The second term $D^{(1)}(\mu^*)$ is also zero, since $D(\mu^* \| \mu)$ has its minimum at $\mu = \mu^*$. Now we rewrite:

$$D^{(2)}(\mu) = \frac{d^2}{d\mu^2} E[\ln M_{\mu^*}(X) - \ln M_\mu(X)] = -\frac{d^2}{d\mu^2} E[\ln M_\mu(X)],$$

which, evaluated at μ^*, resembles the Fisher information. Fisher information is usually defined as $I(\theta) := E\left[(\frac{d}{d\theta}\ln f(X \mid \theta))^2\right]$, but as is well known [14], for exponential families this is equal to $-\frac{d^2}{d\theta^2}E\left[\ln f(X \mid \theta)\right]$, which matches $D^{(2)}(\cdot)$ exactly. Combining this with (8) (Section 3), we obtain:

$$D(M_{\mu^*}\|M_{\hat{\mu}_i}) = \frac{1}{2}{\delta_i}^2/\mathrm{var}_{M_{\mu^*}}(X) + \frac{1}{6}{\delta_i}^3 D^{(3)}(\mu^*) + \frac{1}{24}{\delta_i}^4 D^{(4)}(\mu) \tag{15}$$

We plug this expression back into Equation 14, giving

$$\mathcal{R}_U(n) = \frac{1}{2\mathrm{var}_{M_{\mu^*}}(X)}\sum_{i=0}^{n-1} E_P\left[{\delta_i}^2\right] + R(n), \tag{16}$$

where the remainder term $R(n)$ is given by

$$R(n) = \sum_{i=0}^{n-1} \underset{\hat{\mu}_i \sim M^*}{E}\left[\frac{1}{6}{\delta_i}^3 D^{(3)}(\mu^*) + \frac{1}{24}{\delta_i}^4 D^{(4)}(\mu)\right] \tag{17}$$

where μ and δ_i are random variables depending on $\hat{\mu}_i$ and i. In Lemma 11 we show that $R(n) = O(1)$, giving:

$$\mathcal{R}_U(n) = O(1) + \frac{1}{2\mathrm{var}_{M_{\mu^*}}(X)}\sum_{i=0}^{n-1} E_P\left[(\hat{\mu}_i - \mu^*)^2\right] \tag{18}$$

Note that $\hat{\mu}_i$ is almost the ML estimator. This suggests that each term in the sum of (18) should be almost equal to the variance of the ML estimator, which is $\mathrm{var}X/i$. Because of the slight modification that we made to the estimator, we get a correction term of $O((i+1)^{-2})$ as established in Theorem 5. This theorem gives:

$$\begin{aligned}\sum_{i=0}^{n-1} E_P\left[(\hat{\mu}_i - \mu^*)^2\right] &= \sum_{i=0}^{n-1} O\left((i+1)^{-2}\right) + \mathrm{var}_P(X)\sum_{i=0}^{n-1}(i+1)^{-1} \\ &= O(1) + \mathrm{var}_P(X)\ln n \end{aligned} \tag{19}$$

The combination of (18) and (19) completes the proof. □

8 Building blocks of the proof

The proof of Theorem 1 is based on Lemma 9 and Lemma 11. These Lemmas are stated and proved, respectively, in Section 8.2 and 8.3. The proofs of Theorem 1 and Theorem 2, as well as the proof of both Lemmas, are based on a number of generally useful results about probabilities and expectations of deviations between the average and the mean of a random variable. Below, we first, in Section 8.1, list these deviation-related results.

8.1 Results about Deviations between Average and Mean

Lemma 4 *Let $X, X_1, X_2, \ldots$ be i.i.d. with mean 0. Then we have $E\left[\left(\sum_{i=1}^n X_i\right)^2\right] = n\mathrm{var}(X)$.*

Proof For $n = 0$ the lemma is obviously true. Suppose it is true for some n. For brevity we write $s_n = \sum_{i=1}^n X_i$. Because the mean is zero, we have $E\left[s_n\right] = \sum EX = 0$. Now we compute $E\left[s_{n+1}^2\right] = E\left[\left(s_n + X\right)^2\right] = E\left[s_n^2\right] + 2E\left[s_n\right]EX + E[X^2] = (n+1)\mathrm{var}(X)$. The proof follows by induction. □

Theorem 5 *Let $X, X_1, \ldots$ be i.i.d. random variables, define $\hat{\mu}_n := (n_0 \cdot x_0 + \sum_{i=1}^n X_i)/(n+n_0)$ and $\mu^* = E[X]$. If $\mathrm{var}X < \infty$, then $E\left[(\hat{\mu}_n - \mu^*)^2\right] = O\left((n+1)^{-2}\right) + \mathrm{var}(X)/(n+1)$.*

Proof We define $Y_i := X_i - \mu^*$; this can be seen as a new sequence of i.i.d. random variables with mean 0 and $\mathrm{var}Y = \mathrm{var}X$. We also set $y_0 := x_0 - \mu^*$. Now we have:

$$\begin{aligned}
E\left[(\hat{\mu}_n - \mu^*)^2\right] &= E\left[\left(n_0 \cdot y_0 + \sum_{i=1}^n Y_i\right)^2\right](n+n_0)^{-2} \\
&= E\left[(n_0 \cdot y_0)^2 + 2n_0 \cdot y_0 \sum_{i=1}^n Y_i + \left(\sum_{i=1}^n Y_i\right)^2\right](n+n_0)^{-2} \\
&= O\left((n+1)^{-2}\right) + E\left[\left(\sum_{i=1}^n Y_i\right)^2\right](n+n_0)^{-2} \\
&\overset{(*)}{=} O\left((n+1)^{-2}\right) + n\mathrm{var}(Y)(n+n_0)^{-2} \\
&= O\left((n+1)^{-2}\right) + \mathrm{var}(X)/(n+1),
\end{aligned}$$

where $(*)$ follows by Lemma 4. □

The following theorem is of some independent interest.

Theorem 6 *Suppose $X, X_1, X_2, \ldots$ are i.i.d. with mean 0. If the first $k \in \mathbb{N}$ moments of X exist, then we have Then $E\left[(\sum_{i=1}^n X_i)^k\right] = O\left(n^{\lfloor \frac{k}{2} \rfloor}\right)$.*

Remark It follows as a special case of Theorem 2 of [26] that $E\left[|\sum_{i=1}^n X_i|^k\right] = O(n^{\frac{k}{2}})$ which almost proves this lemma and which would in fact be sufficient for our purposes. We use this lemma instead which has an elementary proof.

Proof We have:

$$E\left[\left(\sum_{i=1}^n X_i\right)^k\right] = E\left[\sum_{i_1=1}^n \cdots \sum_{i_k=1}^n X_{i_1} \cdots X_{i_k}\right] = \sum_{i_1=1}^n \cdots \sum_{i_k=1}^n E[X_{i_1} \cdots X_{i_k}]$$

We define the *frequency sequence* of a term to be the sequence of exponents of the different random variables in the term, in decreasing order. For a frequency sequence $f_1, \ldots, f_m$, we have $\sum_{i=1}^m f_i = k$. Furthermore, using independence of the different random variables, we can rewrite $E[X_{i_1} \cdots X_{i_k}] = \prod_{i=1}^m E[X^{f_i}]$ so the value of each term is determined by its frequency sequence. By computing the number of terms that share a particular frequency sequence, we obtain:

$$E\left[\left(\sum_{i=1}^n X_i\right)^k\right] = \sum_{f_1+\ldots+f_m=k} \binom{n}{m} \binom{k}{f_1, \ldots, f_m} \prod_{i=1}^m E[X^{f_i}]$$

To determine the asymptotic behavior, first observe that the frequency sequence $f_1, \ldots, f_m$ of which the contribution grows fastest in n is the longest sequence, since for that sequence the value of $\binom{n}{m}$ is maximized as $n \to \infty$. However, since the mean is zero, we can discard all sequences with an element 1, because the for those sequences we have $\prod_{i=1}^m E[X^{f_i}] = 0$ so they contribute nothing to the expectation. Under this constraint, we obtain the longest sequence for even k by setting $f_i = 2$ for all $1 \leq i \leq m$; for odd k by setting $f_1 = 3$ and

$f_i = 2$ for all $2 \leq i \leq m$; in both cases we have $m = \lfloor \frac{k}{2} \rfloor$. The number of terms grows as $\binom{n}{m} \leq n^m/m! = O(n^m)$; for $m = \lfloor \frac{k}{2} \rfloor$ we obtain the upper bound $O\left(n^{\lfloor \frac{k}{2} \rfloor}\right)$. The number of frequency sequences is finite and does not depend on n; since the contribution of each one is $O\left(n^{\lfloor \frac{k}{2} \rfloor}\right)$, so must be the sum. □

Theorem 7 *Let $X, X_1, \ldots$ be i.i.d. random variables, define $\hat{\mu}_n := (n_0 \cdot x_0 + \sum_{i=1}^n X_i)/(n+n_0)$ and $\mu^* = E[X]$. If the first k moments of X exist, then $E[(\hat{\mu}_n - \mu^*)^k] = O(n^{-\lceil \frac{k}{2} \rceil})$.*

Proof The proof is similar to the proof for Theorem 5. We define $Y_i := X_i - \mu^*$; this can be seen as a new sequence of i.i.d. random variables with mean 0, and $y_0 := x_0 - \mu^*$. Now we have:

$$\begin{aligned} E\left[(\hat{\mu}_n - \mu^*)^k\right] &= E\left[\left(n_0 \cdot y_0 + \sum_{i=1}^n Y_i\right)^k\right](n+n_0)^{-k} \\ &= O\left(n^{-k}\right) \sum_{p=0}^k \binom{k}{p} (n_0 \cdot y_0)^p E\left[\left(\sum_{i=1}^n Y_i\right)^{k-p}\right] \\ &= O\left(n^{-k}\right) \sum_{p=0}^k \binom{k}{p} (n_0 \cdot y_0)^p \cdot O\left(n^{\lfloor \frac{k-p}{2} \rfloor}\right). \end{aligned}$$

In the last step we used Theorem 6 to bound the expectation. We sum $k+1$ terms of which the term for $p = 0$ grows fastest in n, so the expression is $O(n^{-\lceil \frac{k}{2} \rceil})$ as required. □

Theorem 7 concerns the *expectation* of the deviation of $\hat{\mu}_n$. We also need a bound on the *probability* of large deviations. To do that we have the following separate theorem:

Theorem 8 *Let $X, X_1, \ldots$ be i.i.d. random variables, define $\hat{\mu}_n := (n_0 \cdot x_0 + \sum_{i=1}^n X_i)/(n+n_0)$ and $\mu^* = E[X]$. Let $k \in \{0, 2, 4, \ldots\}$. If the first k moments exists then $P(|\hat{\mu}_n - \mu^*| \geq \delta) = O\left(n^{-\lceil \frac{k}{2} \rceil} \delta^{-k}\right)$.*

Proof

$$\begin{aligned} P(|\hat{\mu}_n - \mu^*| \geq \delta) &= P\left((\hat{\mu}_n - \mu^*)^k \geq \delta^k\right) \\ &\leq E\left[(\hat{\mu}_n - \mu^*)^k\right] \delta^{-k} \quad \text{(by Markov's inequality)} \\ &= O\left(n^{-\frac{k}{2}} \delta^{-k}\right) \qquad \text{(by Theorem 7)} \end{aligned}$$

□

8.2 Lemma 9: Redundancy for Exponential Families

Lemma 9 *Let U be a prequential ML model and $\mathcal{M}$ be an exponential family as in Theorem 1. We have*

$$\mathcal{R}_U(n) = \sum_{i=0}^{n-1} \underset{\hat{\mu}_i \sim P}{E} \left[D(M_{\mu^*} \parallel M_{\hat{\mu}_i})\right].$$

(Here, the notation $\hat{\mu}_i \sim P$ means that we take the expectation with respect to P over data sequences of length i, of which $\hat{\mu}_i$ is a function.)

Proof We have:

$$\arg\inf_{\mu} E_P\left[-\ln M_\mu(X^n)\right] = \arg\inf_{\mu} E_P\left[\ln \frac{M_{\mu^*}(X^n)}{M_\mu(X^n)}\right] = \arg\inf_{\mu} D(M_{\mu^*} \parallel M_\mu)$$

In the last step we used Proposition 10 below. The divergence is minimized when $\mu = \mu^*$ [14], so we find that:

$$\begin{aligned}\mathcal{R}_U(n) &= E_P[-\ln U(X^n)] - E_P[-\ln M_{\mu^*}(X^n)] = E_P\left[\ln \frac{M_{\mu^*}(X^n)}{U(X^n)}\right] \\ &= E_P\left[\sum_{i=0}^{n-1} \ln \frac{M_{\mu^*}(X_i)}{M_{\hat{\mu}_i}(X_i)}\right] = \sum_{i=0}^{n-1} E_P\left[\ln \frac{M_{\mu^*}(X_i)}{M_{\hat{\mu}_i}(X_i)}\right] = \sum_{i=0}^{n-1} \underset{\hat{\mu}_i \sim P}{E}\left[D(M_{\mu^*} \parallel M_{\hat{\mu}_i})\right]. \quad (20)\end{aligned}$$

Here, the last step again follows from Proposition 10. □

Proposition 10 *Let $X \sim P$ with mean μ^*, and let M_μ index an exponential family with sufficient statistic X, so that M_{μ^*} exists. We have:*

$$E_P\left[-\ln \frac{M_{\mu^*}(X)}{M_\theta(X)}\right] = D(M_{\mu^*} \parallel M_\theta)$$

Proof Let $\eta(\cdot)$ denote the function mapping parameters in the mean value parameterization to the natural parameterization. (It is the inverse of the function $\mu(\cdot)$ which was introduced in the discussion of exponential families.) By working out both sides of the equation we find that they both reduce to:

$$\eta(\mu^*)\mu^* + \ln \mathrm{Z}(\eta(\mu^*)) - \eta(\theta)\mu^* - \ln \mathrm{Z}(\eta(\theta)).$$

□

8.3 Lemma 11: Convergence of the sum of the remainder terms

Lemma 11 *Let $R(n)$ be defined as in (17). Then*

$$R(n) = O(1).$$

Proof We omit irrelevant constants and the term for the first outcome, which is well-defined because of our modification of the ML estimator. We abbreviate $\frac{d^k}{d\mu^k} D(M_{\mu^*} \| M_\mu) = D^{(k)}(\mu)$ as in the proof of Theorem 1. First we consider the third order term. We write $E_{\delta_i \sim P}$ to indicate that we take the expectation over data which is distributed according to P, of which δ_i is a function. We use Theorem 7 to bound the expectation of ${\delta_i}^3$; under the condition that the first three moments exist, which is assumed to be the case, we obtain:

$$\sum_{i=1}^{n-1} \underset{\delta_i \sim P}{E}\left[{\delta_i}^3 D^{(3)}(\mu^*)\right] = D^{(3)}(\mu^*) \sum_{i=1}^{n-1} E[\delta_i^3] = D^{(3)}(\mu^*) \sum_{i=1}^{n-1} O(i^{-2}) = O(1).$$

(The constants implicit in the big-ohs are the same across terms.)

The fourth order term is more involved, because $D^{(4)}(\mu)$ is not necessarily constant across terms. To compute it we first distinguish a number of regions in the value space of δ_i: let $\Delta_- = (-\infty, 0)$ and let $\Delta_0 = [0, a)$ for some constant value $a > 0$. If the individual outcomes X are bounded on the right hand side by a value g then we require that $a < g$ and we define

$\Delta_1 = [a, g)$; otherwise we define $\Delta_j = [a + j - 1, a + j)$ for $j \geq 1$. Now we must establish convergence of:

$$\sum_{i=1}^{n-1} \underset{\delta_i \sim P}{E}\left[{\delta_i}^4 D^{(4)}(\mu)\right] \;=\; \sum_{i=1}^{n-1}\sum_{j} P(\delta_i \in \Delta_j) \underset{\delta_i \sim P}{E}\left[{\delta_i}^4 D^{(4)}(\mu) \mid \delta_i \in \Delta_j\right]$$

If we can establish that the sum converges for all regions Δ_j for $j \geq 0$, then we can use a symmetrical argument to establish convergence for Δ_- as well, so it suffices if we restrict ourselves to $j \geq 0$. First we show convergence for Δ_0. In this case, the basic idea is that since the remainder $D^{(4)}(\mu)$ is well-defined over the interval $\mu^* \leq \mu < \mu^* + a$, we can bound it by its extremum on that interval, namely $m := \sup_{\mu \in [\mu^*, \mu^* + a)} \left|D^{(4)}(\mu)\right|$. Now we get:

$$\left|\sum_{i=1}^{n-1} P(\delta_i \in \Delta_0) E\left[{\delta_i}^4 D^{(4)}(\mu) \mid \delta_i \in \Delta_0\right]\right| \;\leq\; \left|\sum_{i=1}^{n-1} 1 \cdot E\left[{\delta_i}^4 \left|D^{(4)}(\mu)\right|\right]\right| \;\leq\; \left|m \sum_i E\left[{\delta_i}^4\right]\right|$$

Using Theorem 7 we find that $E[{\delta_i}^4]$ is $O(i^{-2})$ of which the sum converges. Theorem 7 requires that the first four moments of P exist, but this is guaranteed to be the case: either the outcomes are bounded from both sides, in which case all moments necessarily exist, or the existence of the required moments is part of the condition on the main theorem.

Now we have to distinguish between the unbounded and bounded cases. First we assume that the X are unbounded from above. In this case, we must show convergence of:

$$\left|\sum_{i=1}^{n-1}\sum_{j=1}^{\infty} P(\delta_i \in \Delta_j) E\left[{\delta_i}^4 D^{(4)}(\mu) \mid \delta_i \in \Delta_j\right]\right|$$

We bound this expression from above. The δ_i in the expectation is at most $a + j$. Furthermore $D^{(4)}(\mu) = O(\mu^{k-6})$ by assumption on the main theorem, where $\mu \in [a+j-1, a+j)$. Depending on k, both boundaries could maximize this function, but it is easy to check that in both cases the resulting function is $O(j^{k-6})$. So we get:

$$\ldots \leq \sum_{i=1}^{n-1}\sum_{j=1}^{\infty} P(|\delta_i| \geq a + j - 1)(a + j)^4 O(j^{k-6})$$

Since we know from the condition on the main theorem that the first $k \geq 4$ moments exist, we can apply Theorem 8 to find that $P(|\delta_i| \geq a+j-1) = O(i^{-\lceil \frac{k}{2} \rceil}(a+j-1)^{-k}) = O(i^{-\frac{k}{2}})O(j^{-k})$ (since k has to be even); plugging this into the equation and simplifying we obtain $\sum_i O(i^{-\frac{k}{2}}) \sum_j O(j^{-2})$. For $k \geq 4$ this expression converges.

Now we consider the case where the outcomes are bounded from above by g. This case is more complicated, since now we have made no extra assumptions as to existence of the moments of P. Of course, if the outcomes are bounded from both sides, then all moments necessarily exist, but if the outcomes are unbounded from below this may not be true. We use a trick to remedy this: we map all outcomes into a new domain in such a way that all moments of the transformed variables are guaranteed to exist. Any constant x^- defines a mapping $g(x) := \max\{x^-, x\}$. Furthermore we define the random variables $Y_i := g(X_i)$, the initial outcome $y_0 := g(x_0)$ and the mapped analogues of μ^* and $\hat{\mu}_i$, respectively: $\mu^\dagger$ is defined as the mean of Y under P and $\tilde{\mu}_i := (y_0 \cdot n_0 + \sum_{j=1}^{i} Y_j)/(i + n_0)$. Since $\tilde{\mu}_i \geq \hat{\mu}_i$, we can bound:

$$\begin{aligned}\left|\sum_i P(\delta_i \in \Delta_1) E\left[{\delta_i}^4 D^{(4)}(\mu) \mid \delta_i \in \Delta_1\right]\right| \;&\leq\; \sum_i P(\hat{\mu}_i - \mu^* \geq a) \sup_{\delta_i \in \Delta_1} \left|{\delta_i}^4 D^{(4)}(\mu)\right| \\ &\leq\; \sum_i P(|\tilde{\mu}_i - \mu^\dagger| \geq a + \mu^* - \mu^\dagger) g^4 \sup_{\delta_i \in \Delta_1} \left|D^{(4)}(\mu)\right|\end{aligned}$$

By choosing x^- small enough, we can bring $\mu^\dagger$ and μ^* arbitrarily close together; in particular we can choose x^- such that $a + \mu^* - \mu^\dagger > 0$ so that application of Theorem 8 is safe. It reveals that the summed probability is $O(i^{-\frac{k}{2}})$. Now we bound $D^{(4)}(\mu)$ which is $O((g-\mu)^{-m})$ for some $m \in \mathbb{N}$ by the condition on the main theorem. Here we use that $\mu \leq \hat{\mu}_i$; the latter is maximized if all outcomes equal the bound g, in which case the estimator equals $g - n_0(g - x_0)/(i + n_0) = g - O(i^{-1})$. Putting all of this together, we get $\sup \left|D^{(4)}(\mu)\right| = O((g-\mu)^{-m}) = O(i^m)$; if we plug this into the equation we obtain:

$$\ldots \;\leq\; \sum_i O(i^{-\frac{k}{2}}) g^4 O(i^m) = g^4 \sum_i O(i^{m-\frac{k}{2}})$$

This converges if we choose $k \geq 6m$. We can do this because the construction of the mapping $g(\cdot)$ ensures that all moments exist, and therefore certainly the first $6m$. □

9 Proof of Theorem 2

We use the same conventions as in the proof of Theorem 1. Specifically, we concentrate on the random variables $X_1, X_2, \ldots$ rather than $Z_1, Z_2, \ldots$, which is justified by Equation (14). Let $f(x^n) = -\ln M_{\mu^*}(x^n) - [\inf_{\mu \in \Theta_\mu} - \ln M_\mu(x^n)]$. Within this section, $\hat{\mu}(x^n)$ is defined as the ordinary ML estimator. Note that, if x^n is such that its ML estimate is defined, then $f(x^n) = -\ln M_{\mu^*}(x^n) + \ln M_{\hat{\mu}(x^n)}(x^n)$.

Note $d(n) = E_P[f(X^n)]$. Let $h(x)$ be the carrier of the exponential family under consideration (see Definition 1). Without loss of generality, we assume $h(x) > 0$ for all x in the finite set $\mathcal{X}$. Let $a_n^2 = n^{-1/2}$. We can write

$$\begin{aligned} d(n) \;=\; E_P[f(X^n)] \;=\; & \pi_n \quad E_P[f(X^n) \mid (\mu^* - \hat{\mu}_n)^2 \geq a_n^2] \\ + \; & (1 - \pi_n) \quad E_P[f(X^n) \mid (\mu^* - \hat{\mu}_n)^2 < a_n^2], \end{aligned} \tag{21}$$

where $\pi_n = P((\mu^* - \hat{\mu}_n)^2 \geq a_n^2)$. We determine $d(n)$ by bounding the two terms on the right of (21). We start with the first term. Since X is bounded, all moments of X exists under P, so we can bound π_n using Theorem 8 with $k = 8$ and $\delta = a_n = n^{-1/4}$. (Note that the theorem in turn makes use of Theorem 7 which remains valid when we use $n_0 = 0$.) This gives

$$\pi_n = O(n^{-2}). \tag{22}$$

Note that for all $x^n \in \mathcal{X}^n$, we have

$$0 \leq f(x^n) \leq \sup_{x^n \in \mathcal{X}^n} f(x^n) \leq \sup_{x^n \in \mathcal{X}^n} - \ln M_{\mu^*}(x^n) \leq nC, \tag{23}$$

where C is some constant. Here the first inequality follows because $\hat{\mu}$ maximizes $\ln M_{\hat{\mu}(x^n)}(x^n)$ over μ; the second is immediate; the third follows because we are dealing with discrete data, so that $M_{\hat{\mu}}$ is a probability mass function, and $M_{\hat{\mu}}(x^n)$ must be ≤ 1. The final inequality follows because μ^* is in the interior of the parameter space, so that the natural parameter $\eta(\mu^*)$ is in the interior of the natural parameter space. Because X is bounded and we assumed $h(x) > 0$ for all $x \in \mathcal{X}$, it follows by the definition of exponential families that $\sup_{x \in \mathcal{X}} - \ln M_{\mu^*}(x) < \infty$.

Together (22) and (23) show that the expression on the first line of (21) converges to 0, so that (21) reduces to

$$d(n) = (1 - \pi_n) E_P[f(X^n) \mid (\mu^* - \hat{\mu}_n)^2 < a_n^2] + O(n^{-1}). \tag{24}$$

To evaluate the term inside the expectation further we first Taylor approximate $f(x^n)$ around $\hat{\mu}_n = \hat{\mu}(x^n)$, for given x^n with $(\mu^* - \hat{\mu}_n)^2 < a_n^2 = 1/\sqrt{n}$. We get

$$f(x^n) = -(\mu^* - \hat{\mu}_n)\frac{d}{d\mu} \ln M_{\hat{\mu}_n}(x^n) + n\frac{1}{2}(\mu^* - \hat{\mu}_n)^2 I(\mu_n), \tag{25}$$

where I is the Fisher information (as defined in Section 7) and μ_n lies in between μ^* and $\hat{\mu}$, and depends on the data x^n. Since the first derivative of μ at the ML estimate $\hat{\mu}$ is 0, the first-order term is 0. Therefore $f(x^n) = \frac{1}{2}n(\mu^* - \hat{\mu}_n)^2 I(\mu_n)$, so that

$$\frac{1}{2}ng(n) \inf_{\mu \in [\mu^* - a_n, \mu^* + a_n]} I(\mu) \leq E_P[f(X^n) \mid (\mu^* - \hat{\mu}_n)^2 < a_n^2] \leq \frac{1}{2}ng(n) \sup_{\mu \in [\mu^* - a_n, \mu^* + a_n]} I(\mu),$$

where we abbreviated $g(n) := E_P[(\mu^* - \hat{\mu}_n)^2 \mid (\mu^* - \hat{\mu}_n)^2 < a_n^2]$. Since $I(\mu)$ is smooth and positive, we can Taylor-approximate it as $I(\mu^*) + O(n^{-\frac{1}{4}})$, so we obtain the bound:

$$E_P[f(X^n) \mid (\mu^* - \hat{\mu}_n)^2 < a_n^2] = ng(n)\left(\frac{1}{2}I(\mu^*) + O(n^{-\frac{1}{4}})\right). \tag{26}$$

To evaluate $g(n)$, note that we have

$$E_P[(\mu^* - \hat{\mu}_n)^2] = \pi_n E_P[(\mu^* - \hat{\mu}_n)^2 \mid (\mu^* - \hat{\mu}_n)^2 \geq a_n^2] + (1 - \pi_n)g(n). \tag{27}$$

Using Theorem 5 with $n_0 = 0$ we rewrite the expectation on the left hand side as $\mathrm{var}_P X/n$. Subsequently reordering terms we obtain:

$$g(n) = \frac{(\mathrm{var}_P X)/n - \pi_n E_P[(\mu^* - \hat{\mu}_n)^2 \mid (\mu^* - \hat{\mu}_n)^2 \geq a_n^2]}{1 - \pi_n}. \tag{28}$$

Plugging this into bound (26), and multiplying both sides by $1 - \pi_n$, we get:

$$(1 - \pi_n)E_P[f(X^n) \mid (\mu^* - \hat{\mu}_n)^2 < a_n^2] = \\ \left(\mathrm{var}_P X - n\pi_n E_P[(\mu^* - \hat{\mu}_n)^2 \mid (\mu^* - \hat{\mu}_n)^2 \geq a_n^2]\right)\left(\frac{1}{2}I(\mu^*) + O(n^{-\frac{1}{4}})\right).$$

Since X is bounded, the expectation on the right must lie between 0 and some constant C. Using $\pi_n = O(n^{-2})$ and the fact that $I(\mu^*) = 1/\mathrm{var}_{M_{\mu^*}} X$ (Equation (8)), we get

$$(1 - \pi_n)E_P[f(X^n) \mid (\mu^* - \hat{\mu}_n)^2 < a_n^2] = \frac{1}{2}\frac{\mathrm{var}_P X}{\mathrm{var}_{M_{\mu^*}} X} + O(n^{-\frac{1}{4}}).$$

The result follows if we combine this with (24).

10 Conclusion and Future Work

In this paper we established two theorems about the relative redundancy, defined in Section 2:

1. A particular type of universal code, the *prequential ML code* or *ML plug-in code*, exhibits behavior that we found unexpected. While other important universal codes such as the NML/Shtarkov and Bayesian codes, achieve a regret of $\frac{1}{2}\ln n$, where n is the sample size, the prequential ML code achieves a relative redundancy of $\frac{1}{2}\frac{\mathrm{var}_P X}{\mathrm{var} M_{\mu^*} X}\ln n$. (Sections 2 and 3.)

2. At least for finite sample spaces, the relative redundancy is very close to the expected regret, the difference going to $\frac{1}{2}\frac{\mathrm{var}_P X}{\mathrm{var}_{M_{\mu^*}} X}$ as the sample size increases (Section 4, Theorem 2). In future work, we hope to extend this theorem to general 1-parameter exponential families with arbitrary sample spaces.

Under the heading "Related Work" in Section 2 we list a substantial amount of literature in which the regret for the prequential ML code is proven to grow with $\frac{1}{2}\ln n$. While this may seem to contradict our results, in fact it does not: In those articles, settings are considered where $P \in \mathcal{M}$, and under such circumstances our own findings predict precisely that behavior.

The first result is robust with respect to slight variations in the definition of the prequential ML code: in our framework the so-called "start-up problem" (the unavailability of an ML estimate for the first few outcomes) is resolved by introducing fake initial outcomes. Our framework thus also covers prequential codes that use other point estimators such as the Bayesian MAP and mean estimators defined relative to a large class of reasonable priors. In Section 5.2 we conjecture that no matter what in-model estimator is used, the prequential model cannot yield a relative redundancy of $\frac{1}{2}\ln n$ independently of the variance of the data generating distribution.

11 Acknowledgment

This work was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778. This publication only reflects the authors' views.

References

[1] K. Azoury and M. Warmuth. Relative loss bounds for on-line density estimation with the exponential family of distributions. *Journal of Machine Learning*, 43(3):211–246, 2001. Special issue on Theoretical Advances in On-Line Learning, Game Theory and Boosting, edited by Y. Singer.

[2] A. Barron, J. Rissanen, and B. Yu. The minimum description length principle in coding and modeling. *IEEE Transactions on Information Theory*, 44(6):2743–2760, 1998. Special Commemorative Issue: Information Theory: 1948-1998.

[3] J.O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer Series in Statistics. Springer-Verlag, New York, revised and expanded second edition, 1985.

[4] N. Cesa-Bianchi and G. Lugosi. Worst-case bounds for the logarithmic loss of predictors. *Journal of Machine Learning*, 43(3):247–264, 2001.

[5] B.S. Clarke and A.R. Barron. Information-theoretic asymptotics of Bayes methods. *IEEE Transactions on Information Theory*, IT-36(3):453–471, 1990.

[6] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley Interscience, New York, 1991.

[7] A.P. Dawid. Present position and potential developments: Some personal views, statistical theory, the prequential approach. *Journal of the Royal Statistical Society, Series A*, 147(2):278–292, 1984.

[8] S. de Rooij and P. Grünwald. An empirical study of MDL model selection with infinite parametric complexity. Available at the CoRR arXiv at `http://xxx.lanl.gov/abs/cs.LG/0501028abs.cs.LG/0501028`, 2005.

[9] Y. Freund. Predicting a binary sequence almost as well as the optimal biased coin. In *Proceedings of the Ninth Annual Conference on Computational Learning Theory (COLT' 96)*, 1996.

[10] L. Gerensce'r. Order estimation of stationary Gaussian ARMA processes using Rissanen's complexity. Technical report, Computer and Automation Institute of the Hungaian Academy of Sciences, 1987.

[11] Peter D. Grünwald. MDL tutorial. In Peter D. Grünwald, In Jae Myung, and Mark A. Pitt, editors, *Advances in Minimum Description Length: Theory and Applications*. MIT Press, 2005.

[12] J.A. Hartigan. *Bayes Theory*. Springer-Verlag, New York, 1983.

[13] E.M. Hemerly and M.H.A. Davis. Strong consistency of the pls criterion for order determination of autoregressive processes. *The Annals of Statistics*, 17(2):941–946, 1989.

[14] Robert E. Kass and Paul W. Vos. *Geometrical Foundations of Asymptotic Inference*. Wiley, 1997.

[15] P. Kontkanen, P. Myllymäki, T. Silander, H. Tirri, and P.D. Grünwald. On predictive distributions and Bayesian networks. *Journal of Statistics and Computing*, 10:39–54, 2000.

[16] P. Kontkanen, P. Myllymäki, and H. Tirri. Comparing prequential model selection criteria in supervised learning of mixture models. In T. Jaakkola and T. Richardson, editors, *Proceedings of the Eighth International Conference on Articial Intelligence and Statistics*, pages 233–238. Morgan Kaufman, 2001.

[17] L. Li and B. Yu. Iterated logarithmic expansions of the pathwise code lengths for exponential families. *IEEE Transactions on Information Theory*, 46(7):2683–2689, 2000.

[18] Dharmendra S. Modha and Elias Masry. Prequential and cross-validated regression estimation. *Machine Learning*, 33(1), 1998.

[19] J.A. Rice. *Mathematical Statistics and Data Analysis*. Duxbury Press, 1995.

[20] J. Rissanen. Universal coding, information, prediction and estimation. *IEEE Transactions on Information Theory*, 30:629–636, 1984.

[21] J. Rissanen. Minimum description length principle. In S. Kotz and N.L. Johnson, editors, *Encyclopaedia of Statistical Sciences, Vol. V*, pages 523–527. Wiley, 1986.

[22] J. Rissanen. A predictive least squares principle. *IMA Journal of Mathematical Control and Information*, 3:211–222, 1986.

[23] J. Rissanen. Stochastic complexity and modeling. *The Annals of Statistics*, 14(3):1080–1100, 1986.

[24] J. Rissanen. *Stochastic Complexity in Statistical Inquiry*, volume 15 of *Series in Computer Science*. World Scientific, 1989.

[25] C.Z. Wei. On predictive least squares principles. *The Annals of Statistics*, 20(1):1–42, 1990.

[26] P. Whittle. Bounds for the moments of linear and quadratic forms in independent variables. *Theory of Probability and its Applications*, V(3), 1960.